

CLAIMS

What is claimed is:

1. A storage system comprising:
interface for receiving a client request to access at least one object file and
5 for selecting one of a plurality of control nodes to manage said request, said client request comprising a unique identifier for said object file;
at least one storage center comprising:
a plurality of intelligent storage nodes for storing said object file;
and
10 said control nodes comprising processes for receiving said client request for said object file, for identifying an intelligent storage node that stores said object file from said unique identifier, for accessing said object file in said intelligent storage node identified, and for transferring said object file to said client request.
- 15 2. The storage system as set forth in claim 1, wherein said interface comprises a load balancing switch for selecting one of said control nodes to load balance client requests in said storage system.
- 20 3. The storage system as set forth in claim 1, wherein:
said client request further comprises a client identification and a certificate;
said control nodes further comprise processes for retrieving a client key corresponding to said client identification, for determining whether said client

request is valid based on said client key and said certificate, and for processing said client request if said client request is valid.

4. The storage system as set forth in claim 1, wherein:

5 said unique identifier comprises an object fingerprint that uniquely identifies an object file for said client request; and

said control nodes further comprise processes for identifying said object file in said storage system through said object fingerprint.

10 5. The storage system as set forth in claim 1, wherein said control nodes further comprises a data cache for storing said object file in said control node when transferring said intelligent storage node to said client.

6. The storage system as set forth in claim 1, wherein:

15 said control nodes further comprising:

a reference that identifies at least one intelligent storage node for storage of at least one of said object files;

processes for determining whether said object file is identified on said reference;

20 processes for broadcasting a request for said object file to said intelligent storage nodes using a multi-cast protocol if said object file is not identified on said reference;

said intelligent storage nodes further comprising:

processes for determining whether said object file is stored therein;

processes for broadcasting, to said control nodes, an identification of
at least one of said intelligent storage nodes that stores said object file if
5 said object file is located; and

said control nodes further comprising processes for updating said reference
at said control nodes to include said identification of said object file.

7. The storage system as set forth in claim 1, wherein two of said
10 intelligent storage nodes store an object file.

8. The storage system as set forth in claim 1, further comprising at least
one additional storage center located geographically disparate from said storage
center.

15 9. A method for storing files in a storage system, said method
comprising the steps of:

storing at least one object file in at least one of a plurality of intelligent
storage nodes;

20 receiving a client request to access said object file;

selecting one of a plurality of control nodes to manage said request;

determining, from said control node selected whether said object file is stored in at least one intelligent storage node;

establishing a connection between said control node and said intelligent storage node if said object file is stored in one of said intelligent storage nodes; and

5 transferring, from said intelligent storage node to said client, said object file.

10. The method as set forth in claim 9, wherein the step of selecting one of a plurality of control nodes comprises the step of selecting one of said control nodes to load balance client requests to said storage center.

10

11. The method as set forth in claim 9, further comprising the steps of:
receiving a client request comprising a client identification and a certificate;
retrieving a client key corresponding to said client identification;
determining whether said client request is valid based on said client key and

15

said certificate;

processing said client request if said client request is valid; and

generating an error message if said client request is not valid.

12. The method as set forth in claim 9, wherein the step of determining
20 whether said object file is stored in at least one intelligent storage node comprises the steps of:

receiving a client request comprising an object fingerprint that uniquely identifies an object file for said client request; and

identifying said object file in said storage system through said object fingerprint.

5

13. The method as set forth in claim 9, further comprising the step of caching said object file in said control node when transferring said intelligent storage node to said client.

10

14. The method as set forth in claim 9, wherein the step of determining whether said object file is stored in at least one intelligent storage node comprises the steps of:

storing a reference at said control node that identifies at least one intelligent storage node for at least a subset of said object files;

15

determining whether said object file is identified on said reference;

if said object file is not identified on said reference,

broadcasting a request from said control node for said object file to said intelligent storage nodes using a multi-cast protocol;

20

determining, at said intelligent storage nodes, whether said object file is stored therein;

broadcasting, to said control nodes, an identification of at least one of said intelligent storage nodes that stores said object file if said object file is located; and

5 updating said reference at said control nodes to include said identification of said object file.

15. The method as set forth in claim 14, further comprising the steps of:
establishing a point to point connection between said control node and one of said intelligent storage nodes if said object file is not found from said multi-cast
10 protocol broadcast;

determining whether said object file is located in said intelligent storage node; and

repeating the steps of establishing a point to point connection and determining whether said object file exists with said intelligent storage nodes until
15 said object file is located.

16. The method as set forth in claim 9, further comprising the step of storing an object file in two different intelligent storage nodes.

20 17. The method as set forth in claim 9, further comprising the step of storing an object file in two different intelligent storage nodes, wherein said intelligent storage nodes are located at different geographic locations.

18. A storage cluster comprising:

an interconnect fabric for receiving requests to retrieve an object file, said request comprising a unique file identifier; and

5 a plurality of intelligent storage nodes, wherein an intelligent storage node comprises at least one central processing unit (“CPU”) and a plurality of disk drives for storing a plurality of object files, wherein an intelligent storage node comprises processes for determining whether said object file of said request is located in one of said disk drives, and processes for broadcasting, over said interconnect fabric, an
10 identification of said object file if said object file is located in said intelligent storage node.